Array Tomography Volume Reconstruction

Reconstructing Volumes from Image Stacks Using Fiji

Andrew Olson

version 1.4, June 13, 2012

Contents

List of Tables

1 Download and Install Fiji

What is Fiji? "**F**iji **i**s **j**ust **I**mageJ." ImageJ is a free, extensible, Java-based, open-source image analysis program, based on Wayne Rasband's NIH Image program for Macintosh computers. ImageJ allows for automation of routine tasks using macros, and includes a macro recording feature to make macro creation even easier. In addition, a multitude of user-developed plugins are available for specific image analysis tasks. If you're handy with programming, you can extend these plugins, or write your own. The advantage of Fiji is that it comes with a set of bundled plugins, organized in a coherent menu structure. Fiji is available for Mac, Windows and Linux from: `http://pacific.mpi-cbg.de/wiki/index.php/Downloads`.

Fiji has all of the plugins you need for reconstructing and viewing array tomography volume data, *except* the MultiStackReg plugin, developed by Brad Busse in Stephen Smith's lab. You can download MultiStackReg from Brad's webpage: `http://bradbusse.net/downloads.html`.

The download file is a JAR file (Java archive), and all you need to do is put it in the Fiji plugins folder. For example, on my Windows machine the plugins folder is at: `c:\Program Folders\Fiji\Fiji.app\plugins`.

2 Download the Sample Data Set

For practicing volume reconstruction, you can download a sample data set (0.9 GB, requires SUnet ID and password). This data set is courtesy of Dr. Kristina Micheva, from Stephen Smith's lab.

The data set is from a YFP mouse line in which about 5% of the pyramidal neurons express YFP. The array consists of 31 200 nm coronal sections, including cortex and hippocampus. One set of images is from cortical layer 5, the other is from the CA2 region of the hippocampus. The images are saved in ZVI format (the native image type for Zeiss AxioVision software, with which the images were acquired). Each ZVI file is a multi-channel image from a single z-section. For both cortex and hippocampus, there were two rounds of immunostaining. The following table lists the fluorescence channel and molecular target for the two rounds of immunostaining. (Note that fluorescence channel numbering of the imported images is slightly tricky for Cycle 2.)

Image Data Information				
immunostaining round	**ZVI file channel**	**Fiji titlebar channel**	**fluor**	**molecular target**
Cycle 1	0	0	Alexa 594	PSD95
	1	1	YFP	YFP
	2	2	DAPI	DNA
	3	3	Alexa 647	synapsin
Cycle 2	0	0	Alexa 594	gephyrin
	2	1	DAPI	DNA
	3	2	Alexa 647	GAD

Table 1: Image Data Information. Note that the fluorescence channel numbering is annoyingly tricky. In the ZVI file metadata, the fluorescence channels have a fixed numbering scheme (bearing some relation to the filter cubes in the microscope). The Bio-Formats Importer plugin does not use this information, but instead numbers the channels sequentially. To make things more confusing, the channel numbering starts from 0 in the image stack titlebar, but starts from 1 in the image stack status line. Because the titlebar is more prominent, we use the titlebar (zero-based) numbering scheme.*

*It should be noted that the plugin software authors are in good company with their counterparts at Zeiss. While the ZVI file channel numbering scheme is zero-based, the scheme the user sees in the AxioVision software interface is one-based.

The diagram below shows how the data is organized on the file server. We suggest that you keep this same directory structure when you copy the data to your computer.

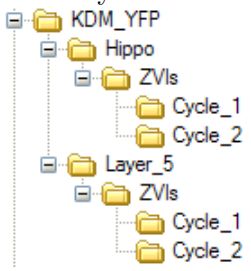

3 Open Sample Data in Fiji

The next step is to open the sample data set using Fiji. First start Fiji, then use the menus to select: `Plugins/LOCI/Bio-Formats Importer`. Navigate to the sample data directory, and select the first ZVI file for Layer 5, Cycle 1 (`YFPH syn-psd 15-5065_1000000.zvi`). You will then be presented with a set of options for importing the images into ImageJ, as shown in the following screen shot:

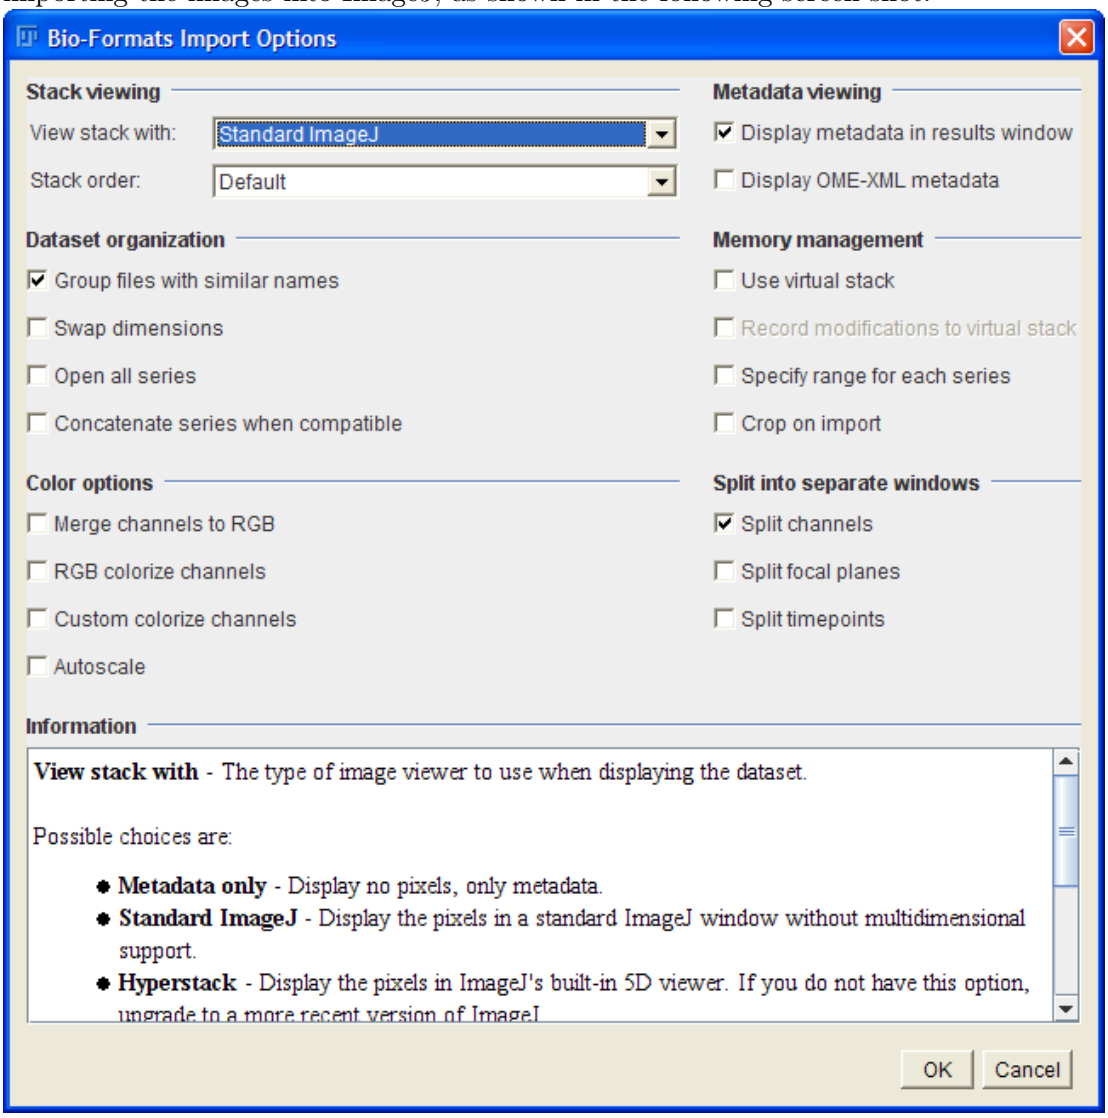

As shown above, you should make the following selections:

- `View stack with: Standard ImageJ`
- `Display metadata in results window` (optional: displays all of the metadata from the ZVI file in a separate window)
- `Group files with similar names` (Creates image stacks from all 31 ZVI files in the sample data set. Note: uncheck this option if your ZVI file contains an entire z-stack, i. e., if you have a single ZVI file with all of your z-sections, rather than separate ZVI files for each z-section.)
- `Split channels` (opens each channel as a separate stack)
- Click `OK` to continue.

The Bio-Formats Importer plugin then analyzes the ZVI image files to determine how to sort the images into z-stacks. Depending on how fast your computer is, you may see the status line of the Fiji window look something like this:

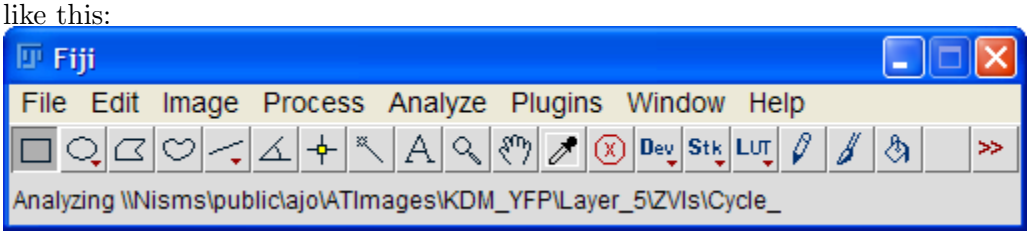

If you had the "`Group files`" option checked, you'll see a window pop up on your screen to tell you how the Bio-Formats Importer plugin has analyzed the image file names:

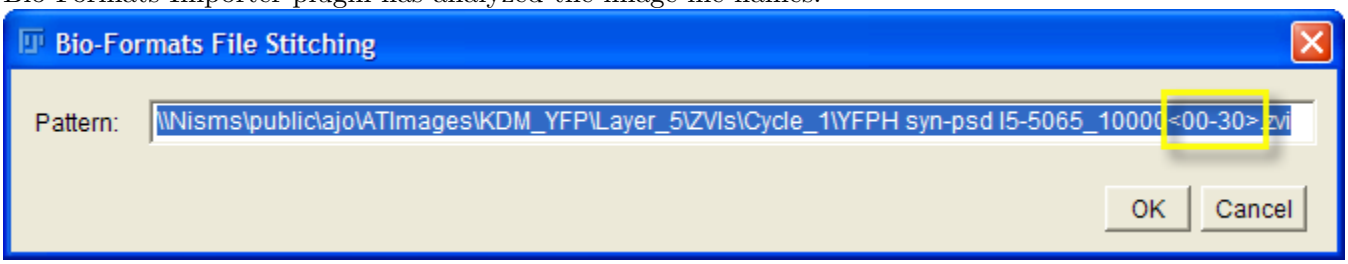

The highlighted portion (in the angled brackets) indicates that the plugin has analyzed the ZVI file directory and found all 31 slices, and will arrange them in order in a stack. Click `OK` to continue.

Next the plugin reads the image metadata (again the Fiji window status line reflects this):

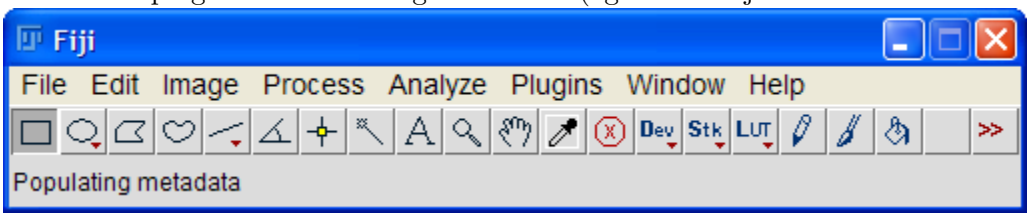

When complete, a results window pops up, showing you the extensive metadata from the ZVI file. The metadata is useful information about how the image was acquired, including exposure time, filter sets used, and scale factors (in μm/pixel). You can resize the window (and the width of the data columns) by clicking and dragging with the mouse. You can also save this information as a plain text file.

Finally, the plugin reads through all of the ZVI files, and opens a separate image stack for each fluorescence channel in the ZVI file (because we selected the `Split channels` option previously). Note that the plugin does *not* use the ZVI file metadata to number the fluorescence channels. It simply numbers them sequentially. This matters for the Cycle 2 sample data, where the ZVI file has no channel 1 image, so the channel numbers assigned

by the Bio-Formats Importer plugin will differ from the ZVI file channel numbers (refer to Table 1, above, for full information on the image data and channel numbering). Again, depending on how fast your computer is, you may see a progress bar in the status line of the main Fiji window, showing you how this is proceeding:

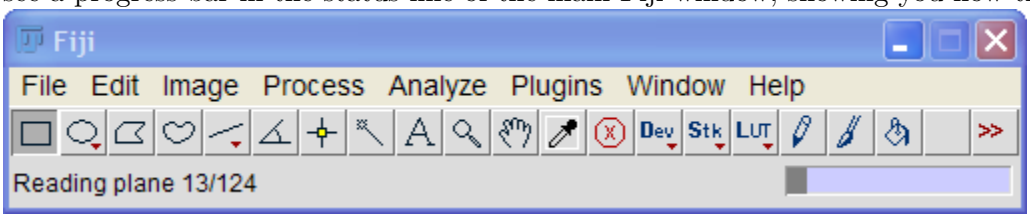

When the process is completed, you will have four image stacks open, one for each fluorescence channel. You can tell which image stack is which by examining the title bar. For example, the title bar below:

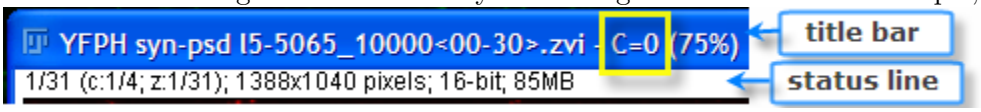

identifies this image stack as channel 0 (synapsin). The status line indicates that you are looking at slice 1 of 31. The status line also has additional information about the image (including a potentially confusing reference to channels using a different numbering scheme). Use the slider at the bottom of the window to navigate through the image stack. You will see the z-position updated in the status line as you do.

4 Image Processing: Subtract Background, Normalize . . .

The additional pre-processing steps described below are optional. In general, it is of course the best practice to optimize your acquisition settings when you initially scan your arrays, so that your images have low background levels and utilize the full linear dynamic range of your microscope's digital camera. But you knew that already. (If you didn't, please see [1, 2] for excellent, practical discussions of microscopy data acquisition.)

4.1 Image Processing: Subtract Background

But this is biology, and you may find that some pre-processing of your images is warranted. For example, you may want to use the "rolling ball" background subtraction algorithm to lower the background gray level. Use the menus to select `Process/Subtract Background ...`. The default value of 50.0 pixels for the diameter of the rolling ball typically works well. The ImageJ documentation has an excellent description of how background subtraction is implemented, and how to use the available parameter and settings.

4.2 Image Processing: Normalization

You may also want to normalize the images in each image stack. Normalization stretches out the pixel values to match the available dynamic range of the image type. Normalization can be used to correct for variation in overall fluorescence between sections (due to differences in section thickness, for example). You can learn about the details of how this is done by consulting the ImageJ documentation. Use the menus to select `Process/Enhance contrast`.

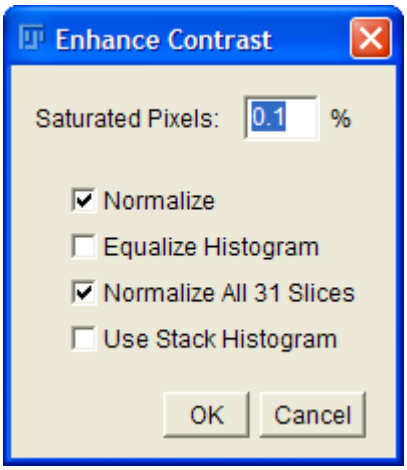

As shown above, make the following selections:

- choose the percentage of pixels in which to allow saturation (suggested starting points are 0.5% for the DAPI channel, and 0.1% for the other channels)
- select `Normalize`
- select `Normalize all 31 slices` (Each slice is independently normalized. This is the typical selection, and will compensate for conditions such as variations in section thickness. If you select `Use stack histogram`, then a single histogram is constructed from the entire z-stack, and each slice is normalized using this single "global" stack histogram.)

Normalization may cause problems with irregular stains. For example, sometimes in the YFP channel, one cell is particularly bright, but is not present in all sections. In this case, normalization will suppress the overall gray level of sections with the bright cell, with an undesirable dimming effect for those sections when viewed in 3D.

4.3 Image Processing: Save Your Work

Finally, after pre-processing, you should save each image stack. Before doing this, it's a good idea to check `File/Image Properties` to make sure that the scale units and the x-, y-, and z-scale factors have been set correctly. For `Unit of Length` fill in μm (type "`um`"). The correct `Pixel Width` and `Pixel Height` can be obtained from the metadata of the ZVI file ("`Scale Factor for X`" and "`Scale Factor for Y`", respectively; see p. 3 for instructions on accessing the metadata). `Voxel Depth` is your section thickness, in μm. You can check the `Global` button in the `Image Properties` dialog, so that these settings will automatically be applied to every image that you save during your Fiji session. Once the `Image Properties` are set, use the menus to select `File/Save As/tiff`. To organize the data, we suggest creating a fresh directory for each step of the process. This will keep your data well-organized, and the directory structure for your images will reflect the actual processing steps that the images have undergone. For example, at this stage, you could create a folder called `RawBackSubNorm`. "`Raw`" because the images have not yet been aligned. The rest of the descriptors are self-explanatory. Each channel will be saved as a separate tiff stack. For naming these files, we suggest using a name that (at a minimum) will identify the molecular target and the immunostaining round. The diagram below shows our resulting files after normalizing all of the Layer 5 data:

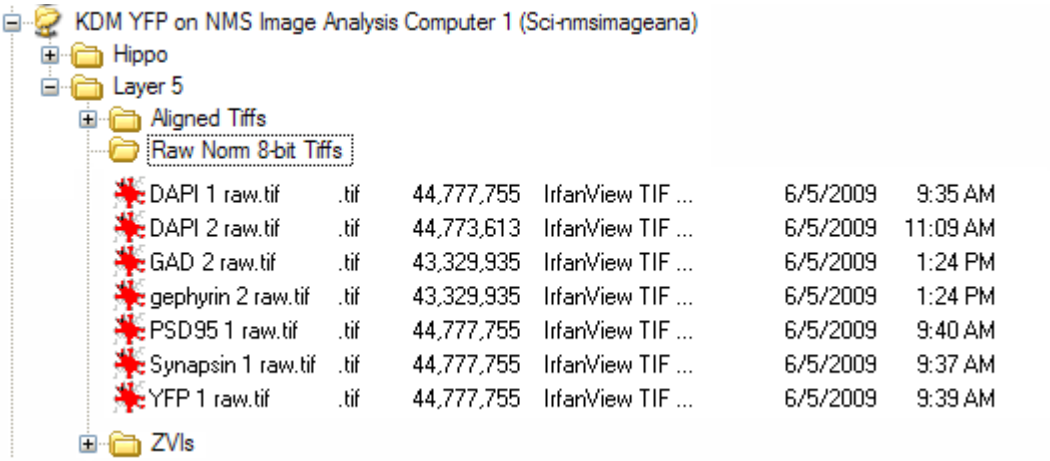

5 Image Alignment

Now we are ready to do the image alignment, using the MultiStackReg plugin with the normalized, background-subtracted images we just created as the starting point. Here's an overview of how it works:

1. Start by aligning with the DAPI image stack from the first round of immunostaining (Cycle 1). The transformation matrices from this step are saved to a file for aligning the image stacks from the other fluorescence channels.

2. Apply the saved transformation matrices from the previous step to the other Cycle 1 image stacks.

3. Align the DAPI image stack from the second round of immunostaining (Cycle 2) to the raw (unaligned) Cycle 1 DAPI image stack. The transformation matrices from this step are saved to a file for use in the next step.

4. Apply the saved transformation matrices from the previous step to the other Cycle 2 image stacks.

5. Finally, apply the Cycle 1 transformation matrices to all of the Cycle 2 image stacks.

That's the basic procedure. More detailed instructions for each step are provided in the following sections.

5.1 Image Alignment: Cycle 1, DAPI Channel

The MultiStackReg plugin works by performing cross-correlations between pairs of adjacent images in the stack. It applies a transform to one image and adjusts the parameters of the transform to maximize the cross-correlation between the two adjacent images. For the sample data set, the DAPI channel the best starting point for alignment. Nuclei are distributed fairly uniformly over the field, but the staining is sparse enough so that spurious correlations are avoided. Using a synaptic marker (e.g., synapsin or PSD95) for the initial alignment is not as efficient, because of the higher likelihood of spurious correlations. As we will see below, these channels can be used to refine the initial alignment from the DAPI channel.

5.1.1 Position the Image Stack

The alignment begins at the current location (i.e., z-position) of the image stack. Use the slider at the bottom of the DAPI image stack window to navigate back and forth through the stack. If the ribbon of sections was curved on the slide, you will notice that the images rotate relative to one another as you navigate through the stack.

The alignment process will correct this by rotating the images, but as a consequence, triangular black regions will be added at the edges of the image. In order to minimize the area of these black regions, start the alignment at the center of the rotation. For the sample data set, this will mean starting at the center of the stack, since the curvature was uniform throughout the ribbon. If your ribbon started out straight and then curved, you'll want to start at the center of the curved section.

Suggestion: if you want to see how big an effect curvature of the ribbon has on alignment, try aligning the images twice. For the first trial, start at the center of the image stack. For the second trial, start at the beginning of the image stack. Save the resulting aligned images in separate directories, and then compare the size of the black triangular regions at the bottom end of each stack.

5.1.2 Start the Alignment

Once you have the DAPI image stack properly positioned, use the menus to select `Plugins/Registration/MultiStackReg`.

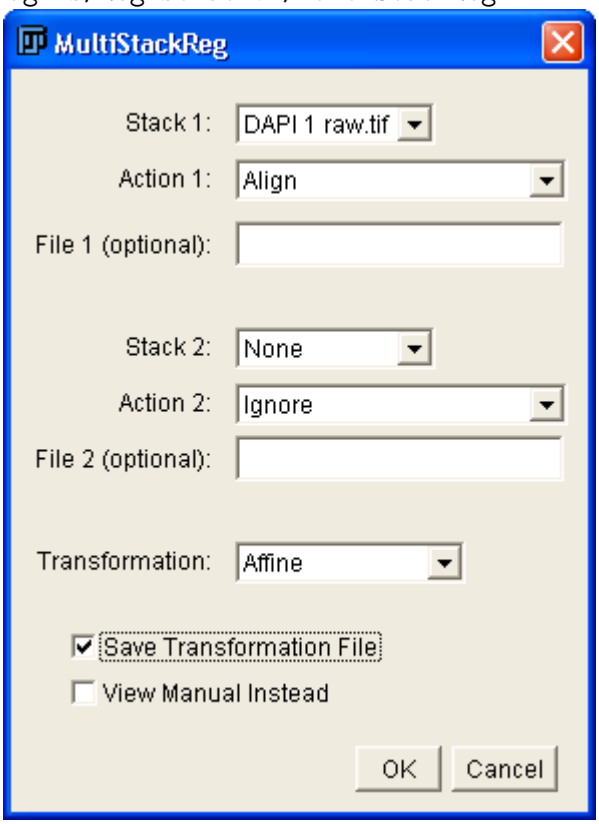

As shown above, select the following MultiStackReg options:

- For `Stack 1`, choose `DAPI 1 raw`.
- For `Action 1`, choose `Align`.
- Ignore `File 1` and the Stack 2 settings, as they are not needed at this point.
- For `Transformation`, choose `Affine`.
- Tick the checkbox for `Save Transformation File` (you will be prompted for a filename later).
- Click `OK` to start the alignment process.

You will first be prompted for a filename to save the transformation matrices. You will use this file in the next step to align the remaining "Cycle 1" image stacks. We suggest that you create a new directory, "`Aligned tiffs`" in which to save the transformation matrices and aligned image stacks. Give the transformation file a good descriptive name, such as: "`DAPI 1 Affine Matrices.txt`".

The alignment process may take awhile to complete, so be patient. For the impatient, there are two visual cues to confirm that MultiStackReg is working:

1. The progress bar (background only) will appear at the bottom right of the main Fiji window while each image is aligned. You'll see it flash off occasionally between images.

2. The slider at the bottom of the DAPI image stack window will move to indicate the image that is currently being aligned. When MultiStackReg is complete, the slider will return to the original starting position.

When the alignment is complete, save the aligned DAPI image stack in the "`Aligned tiffs`" directory. Now you're ready to align the rest of the channels from "Cycle 1".

5.2 Image Alignment: Cycle 1, Additional Channels

Aligning the rest of the channels from Cycle 1 is straightforward. You simply apply the saved transformation matrices from the DAPI 1 image stack to each of the additional Cycle 1 stacks, one-by-one. Again, use the menus to select `Plugins/Registration/MultiStackReg`.

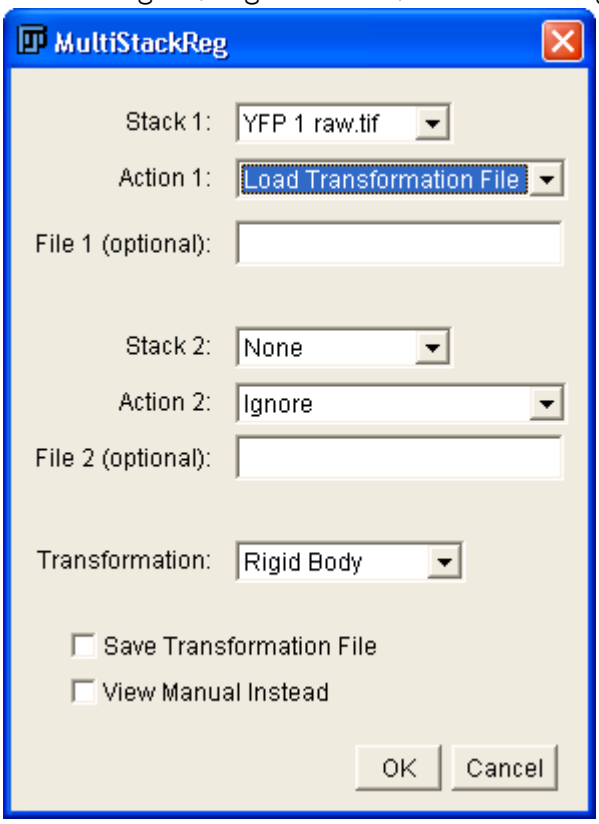

As shown above, select the following MultiStackReg options:

- For `Stack 1`, choose (e.g.) `YFP 1 raw`.

- For `Action 1`, choose `Load Transformation File`.

- Do not tick the checkbox for `Save Transformation File` (not necessary, it's already saved and this step won't change it).
- Ignore the remaining settings, as they are not needed at this point.
- Click `OK` to start the alignment process.

This alignment should proceed rapidly, because the computer is only applying existing transformations, not calculating new ones. When complete, you should again save the aligned image stack in the "`Aligned tiffs`" directory.

Repeat this step for all of the remaining "Cycle 1" images.

5.3 Image Alignment: Cycle 2, DAPI Channel

Now you're ready to begin the alignment of the "Cycle 2" image stacks. This will be a two-step process.

- You will start by aligning the raw DAPI 2 image stack to the raw (unaligned) DAPI 1 image stack. Save the resulting transformation matrices to a file, and apply it to the remaining "Cycle 2" image stacks.
- Complete the "Cycle 2" alignment by applying the "Cycle 1" transformation matrices to the "Cycle 2" stacks.

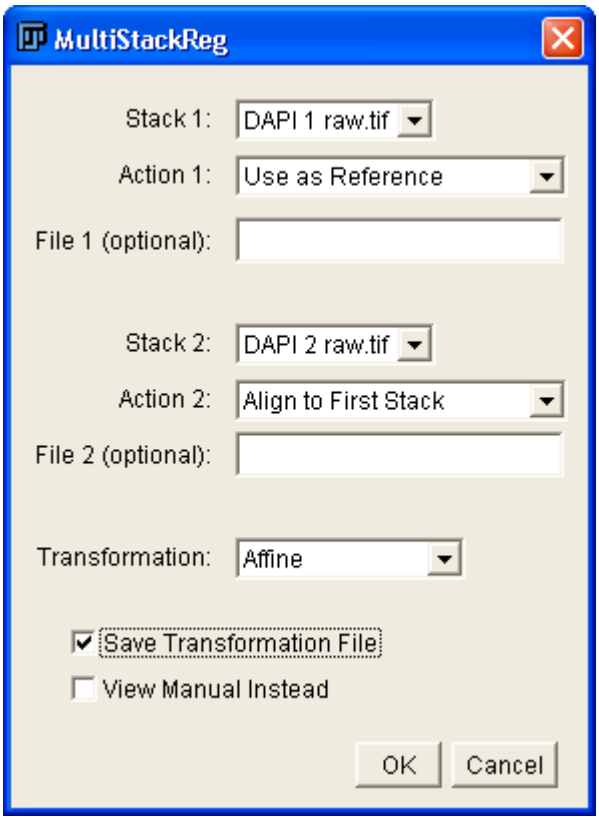

First open the raw (unaligned) DAPI image stacks. Then, as shown above, select the following MultiStackReg options:

- For `Stack 1`, choose `DAPI 1 raw`.
- For `Action 1`, choose `Use As Reference`.

- Ignore `File 1` setting.
- For `Stack 2`, choose `DAPI 2 raw`.
- For `Action 2`, choose `Align to First Stack`.
- For `Transformation`, choose `Affine`.
- Tick the checkbox for `Save Transformation File` (you will be prompted for a filename later).
- Click `OK` to start the alignment process.

Use a good, descriptive name for saving the transformation file, such as: `DAPI 2 raw to DAPI 1 raw Affine Matrices.txt`.

In this case the starting position of the image stack doesn't matter, because the alignment is done pair-wise at each z-position between the DAPI 1 raw image stack and the DAPI 2 raw image stack. If you watch the image stack windows, you'll see the slider at the bottom of the window marching through the two stacks from top to bottom. When the alignment is complete, the slider for the target stack moves to the center of the stack, but for the newly aligned stack, the slider remains at the last image in the stack. You must move this slider to another image in the stack in order to fully apply the transform to the final image.

Now you have the DAPI 2 raw image aligned to the DAPI 1 raw image. In order to complete the alignment of the DAPI 2 stack, apply the cycle 1 transformation matrices (i.e., "`DAPI 1 Affine Matrices.txt`") using the Load Transformation File action in MultiStackReg.

Finally, save the fully aligned DAPI 2 image stack in the "`Aligned tiffs`" directory.

A good way to do a quick quality-control check on the alignment between the two DAPI image stacks is to use the "Color Merge" feature of ImageJ. This allows you to superimpose the image stacks, displaying one with red and the other green. Where the two images overlap (i.e., align perfectly), you'll see yellow. Where the two images don't overlap (i.e., there is a problem with alignment), you'll see the red and green colors of the individual stacks. Open the two aligned image stacks (`DAPI 1` and `DAPI 2`), and then use the menus to select `Image/Color/Merge Channels`.

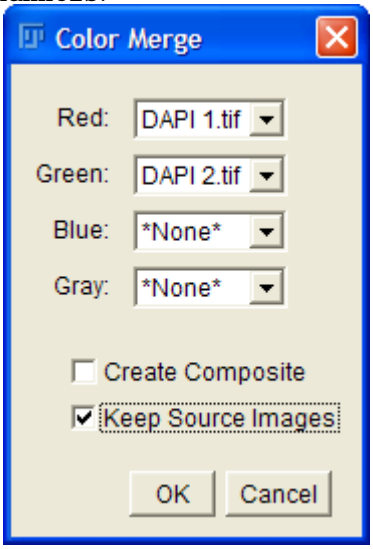

As shown above, make the following selections in the `Color Merge` dialog:

- For `Red`, choose `DAPI 1`.
- For `Green`, choose `DAPI 2`.

- For `Blue`, choose `*None*`.
- For `Gray`, choose `*None*`.
- Un-tick the checkbox for `Create Composite`.
- Tick the checkbox for `Keep Source Images`.
- Click `OK`.

Use the slider at the bottom of the newly-created RGB window to navigate through the color-merged image stack. Remember that yellow indicates good alignment, and that problems are flagged by the red and green areas.

You will notice that there are significant problems with a some of the z-sections. Closer examination of the DAPI 2 image stack will explain what's going on. In z-sections 3, 22 and 28 there are "extra" nuclei (or other debris) present in the image. This happened because the antibody stripping process can sometimes cause the sections themselves to come loose. Parts of one section may land atop another, with the results that you see. The Smith lab is working on developing more resilient methods for attaching the sections to the cover slip using UV-curing adhesives. Consult with them if you are having difficulty with sections coming loose during the antibody stripping process.

Meanwhile, the simplest way to rescue the Cycle 2 image data is to erase (black out) the problematic sections and use the remaining data in the stack. To black out a section, use the rectangular ROI (region of interest) tool (highlighted in yellow, below) to select the entire section:

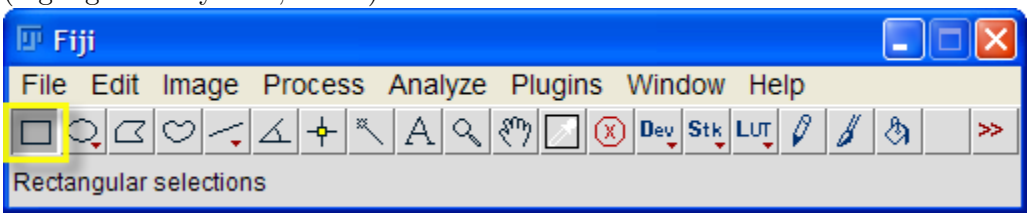

Then use the menus to select `Edit/Fill` (or type `Ctrl+F`). (Note that, if needed, you can set the fill color to black via `Edit/Options/Colors`.) If the `Process Stack` dialog pops up:

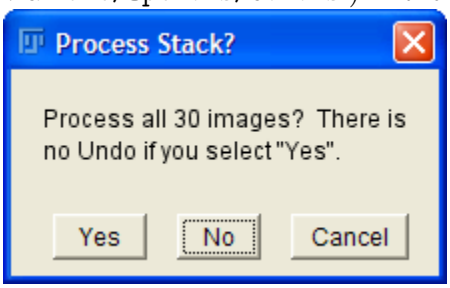

(on my machine it does if you select `Edit/Fill`, but does not if you type `Ctrl+F`), click `No` in order to black out only the selected z-section, *not* the entire stack. Repeat this for each affected z-section, then save the processed image stack with a descriptive name, such as: `DAPI 2 blanked z3 z22 z28.tif`.

5.4 Image Alignment: Cycle 2, Additional Channels

Finally, you'll need to apply two rounds of transformation matrices to the remaining cycle 2 images:

- `DAPI 2 raw to DAPI 1 raw Affine Matrices.txt`, followed by
- `DAPI 1 Affine Matrices.txt`.

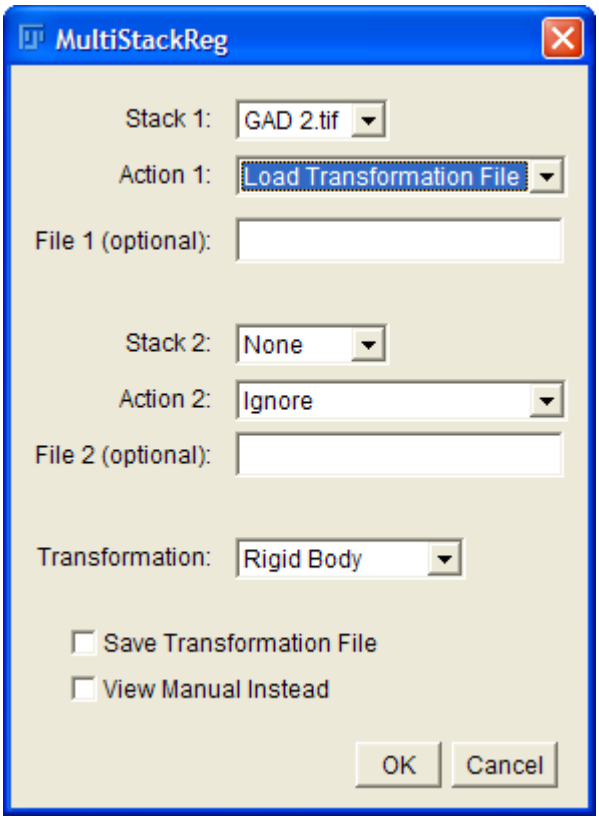

As shown above, select the following MultiStackReg options:

- For `Stack 1`, choose (e.g.) `GAD 2`.
- For `Action 1`, choose `Load Transformation File`.
- Do not tick the checkbox for `Save Transformation File` (not necessary, it's already saved and this step won't change it).
- Ignore the remaining settings, as they are not needed at this point.
- Click `OK` to start the alignment process.
- Select the first transformation file, then repeat the process for the second transformation file.

Remember to save your aligned image stacks!

6 Viewing the Aligned Stacks in 3D

Finally, the payoff: seeing the painstakingly reconstructed image data in glorious 3D! This section will provide a brief introduction to Fiji's 3D Viewer plugin. For more detailed information on the 3D Viewer, visit its webpage: http://132.187.25.13/home/?category=Download&page=Viewer3D.

6.1 Loading Image Stacks in the 3D Viewer

Open the image stacks that you wish to view, then use the menus to select `Plugins/3D Viewer`. The viewer window opens, along with its `Add...` dialog box (shown below),

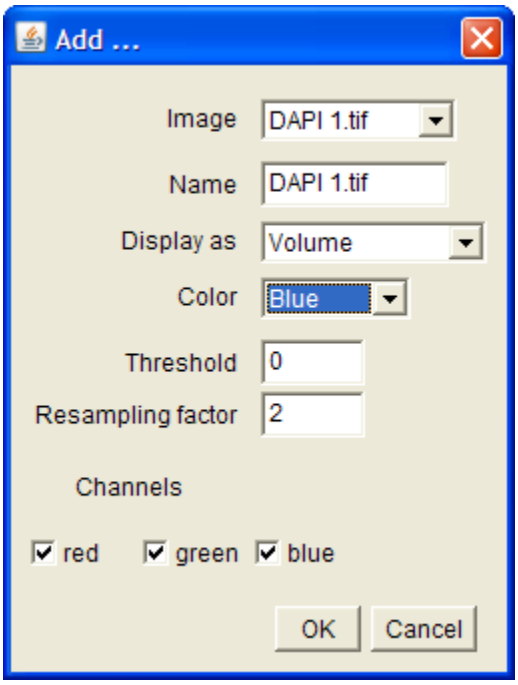

which you use for adding image stacks to the viewer, one by one. Using the drop-down selections in the dialog, choose the image stack to display (`Image`), and the pseudocolor for displaying the stack (`Color`). In the example above, we're about to load the `DAPI1.tif` image stack, to be displayed in blue. For `Display as`, use the default choice of `Volume`. Leave the `Threshold` at the default setting of 0. For starters, leave the `Resampling factor` at the default setting of 2. (This setting allows faster 3D display at the expense of resolution. For full resolution display, change `Resampling factor` to 1.) The `Channels` check boxes allow you to select/deselect individual channels from a multi-channel image stack. Click on `OK`, and wait while the image stack loads (note the progress bar in the Fiji main window status line).

The DAPI image stack alone is pretty boring, and we've got three primary colors, so let's add a couple more image stacks. Use the 3D Viewer menus (note: *not* the Fiji menus this time!) to select `File/Add content`. This brings up the `Add...` dialog again. Display the YFP1.tif stack using green. [Lather, rinse, and ;-)] Repeat the process to add either the Synapsin1.tif or the PSD951.tif stack in red.

6.2 3D Viewer Display Controls

The following table describes the "freehand" controls that you can use to manipulate the view (i.e., rotate, translate, zoom) of your image data in the 3D Viewer. (Note: "hand" and "glass" refer to selections from the Fiji toolbar, the scroll tool (hand) and the magnifying glass tool (glass)).

function	mouse	keyboard
rotate	hand + L button	hand (or glass) + arrow keys
translate	hand + shift + L button	hand (or glass) + shift + arrow keys
zoom	glass + L button (or mouse wheel)	hand (or glass) + page up/page down keys

With the arrow keys, rotation is restricted about either the x- or y-axis. With the mouse, you can achieve rotation at any angle [*except* the one you want, it sometimes seems ;-)]. You can always use the 3D Viewer menus to select `View/Reset view` to get back to the starting point.

6.3 Tips and Tricks

Here's a potential *gotcha!*: note that by clicking the mouse you can select an individual image stack. It will be highlighted with a red seletion frame. If you start dragging the mouse at this point, you will rotate *only* the selected image stack, leaving the others behind, which is probably not what you want. Click somewhere in the viewer background to deselect. If, perchance, you do end up with individual stacks out of register with the others, you can select them (with the mouse, or via the `Select` menu), and then delete them from the 3D Viewer using `File/Delete`. This is also the procedure to follow for changing the image data you have displayed.

Displaying array tomography data sets with the 3D Viewer can easily push your computer's video card to its limits. If the 3D Viewer display freezes (usually temporary), this is probably what's going on.

You can also use the 3D Viewer to make movies of your data. The simplest way is select `View/Start 360 deg recording`. This will make a movie in which your image stacks rotate 360 degrees around the y-axis. It may take awhile, but be patient. When complete, a movie window will open, in which you can replay the rotation, and from which you can save the movie (as an AVI file). The viewer also has a freehand recording feature (start and stop commands available from the 3D Viewer `View` menu. For smoothest results, you'll probably want to use keyboard navigation.

Again, for more detailed information, consult the 3D viewer webpage.

7 Conclusion

I hope you found this introduction useful. For further assistance with array tomography reconstruction and analysis, please contact the Neuroscience Microscopy Service.

8 Acknowledgements

Many thanks to Prof. Stephen Smith and Dr. Kristina Micheva for their time and assistance teaching me these techniques, and for their helpful comments. Any errors in this tutorial are my own. Suggestions for improving this documentation are always welcome.

References

[1] Alison J North. Seeing is believing? a beginners' guide to practical pitfalls in image acquisition. *J Cell Biol*, 172(1):9–18, January 2006.

[2] Helen Pearson. The good, the bad and the ugly. *Nature*, 447(7141):138–40, May 2007.